arXiv:1407.4824v1 [q-bio.MN] 17 Jul 2014

Fair Evaluation of Global Network Aligners

Joseph Crawford,[*†‡] Yihan Sun,[*§] and Tijana Milenković[*†‡¶]

July 21, 2014

Abstract

Analogous to genomic sequence alignment, biological network alignment identifies conserved regions between networks of different species. Then, function can be transferred from well- to poorly-annotated species between aligned network regions. Network alignment encompasses two algorithmic components: node cost function (NCF), which measures similarities between nodes in different networks, and alignment strategy (AS), which uses these similarities to rapidly identify high-scoring alignments. Different methods use both different NCFs and different ASs. Thus, it is unclear whether the superiority of a method comes from its NCF, its AS, or both. We already showed on state-of-the-art methods, MI-GRAAL and IsoRankN, that combining NCF of one method and AS of another method *can* lead to a new superior method. Here, we evaluate MI-GRAAL against a newer approach, GHOST, by mixing-and-matching the methods' NCFs and ASs to potentially further improve alignment quality. While doing so, we approach several important questions that have not been asked systematically thus far. First, we ask how much of the node similarity information within NCF should come from protein sequence data compared to network topology data. Existing methods determine this parameter more-less arbitrarily, which could significantly affect the resulting alignment(s). Second, when topological information is used in NCF, we ask how large the size of the neighborhoods of the compared nodes should be. Existing methods assume that the larger the neighborhood size, the better.

We find that MI-GRAAL's NCF is superior to GHOST's NCF, while the performance of the methods' ASs is data-dependent. Thus, for data on which GHOST's AS is superior to MI-GRAAL's AS, the combination of MI-GRAAL's NCF and GHOST's AS represents a new superior method. Also, we find that which amount of sequence information is used within NCF does not affect alignment quality, while the inclusion of topological information is crucial for producing good alignments. Finally, we find that larger neighborhood sizes are preferred, but often, it is the second largest size that is superior. Using this size instead of the largest one would decrease computational complexity.

Taken together, our results lead to several general recommendations for a fair evaluation of network alignment methods.

1 Background

1.1 Motivation and related work

Analogous to sequence alignment, which finds regions of similarity that are a likely consequence of functional or evolutionary relationships between the sequences, network (or graph) alignment finds regions of topological and functional similarity between networks of different species [35]. Then,

[*]Department of Computer Science and Engineering, University of Notre Dame
[†]Interdisciplinary Center for Network Science and Applications (iCeNSA), University of Notre Dame
[‡]ECK Institute for Global Health, University of Notre Dame
[§]Department of Computer Science and Technology, Tsinghua University
[¶]Corresponding Author (E-mail: tmilenko@nd.edu)

functional (e.g., aging-related [28, 9, 10]) knowledge can be transferred between species across conserved (aligned) network regions. Thus, just as sequence alignment, network alignment can be used for establishing from biological network data orthologous relationships between different proteins or phylogenetic relationships between different species [18, 26, 19]. Also, it can be used to semantically match entities in different ontologies [20].

Network alignment can be performed locally and globally. Local network alignment (LNA) aims to optimize similarity between local regions of different networks [15, 36, 11, 17, 2, 21, 3, 29, 5]. As such, LNA often leads to many-to-many node mapping between different networks. However, LNA is generally unable to find large conserved subgraphs. Thus, methods for global network alignment (GNA) have been proposed, which aim to optimize global similarity between different networks and can thus find large conserved subgraphs [38, 12, 37, 42, 22, 18, 26, 19, 32, 31, 30, 13, 16, 8, 22, 28, 9, 34, 6]. Unlike LNA, GNA typically results in one-to-one node mapping between different networks. In this study, we focus on GNA due to its recent popularity, but all concepts and ideas can be applied to LNA as well.

More formally, we define GNA as a one-to-one mapping between nodes of two networks that aligns the networks well with respect to a desired topological or functional criterion. GNA is a computationally hard problem to solve due to the underlying subgraph isomorphism problem [41]. This is an NP-complete problem that asks whether a network exists as an exact subgraph of a larger network. GNA is a more general problem which aims to fit well two networks when one network is not necessarily an exact subgraph of another network. Since GNA is computationally hard, heuristic methods need to be sought. Existing GNA heuristic algorithms typically achieve an alignment via two algorithmic components: node cost function (NCF) and alignment strategy (AS). NCF captures pairwise costs (or equivalently, similarities) of aligning nodes in different networks, and AS uses these costs to identify a good-quality alignment out of all possible alignments with respect to some topological or biological alignment quality measure.

Different existing GNA methods generally use both different NCF and AS, so it is unclear whether the superiority of a method comes from its NCF, AS, or both. For this reason, in our recent study [28, 9], we combined NCFs and ASs of MI-GRAAL [19] and IsoRankN [22], two state-of-the-art methods at the time, as a proof of concept that it is important to fairly evaluate the contribution of each component to alignment quality. In the process, we showed that NCF of MI-GRAAL is superior to that of IsoRankN, and importantly, we proposed the combination of MI-GRAAL's NCF and IsoRankN's AS as a new superior method for multiple GNA, i.e., for GNA of more than two networks at a time [28, 9].

In the meanwhile, a new state-of-the-art method has appeared, called GHOST [32] (along with some other methods [6]). Thus, in this study, we fairly evaluate MI-GRAAL against GHOST by mixing and matching their NCF and AS. At the same time, we ask several additional important questions regarding the choice of appropriate GNA parameters, which have surprisingly been neglected thus far.

1.2 Our approach and contributions

MI-GRAAL [19] and GHOST [32] are two state-of-the-art global network aligners that injectively map nodes between two networks in a way that preserves topologically or functionally conserved network regions. The two methods are conceptually similar, in the sense that their NCFs assume two nodes from different networks to be similar if their topological neighborhoods are similar. However, the mathematical and implementation details of the two NCFs are different. The same holds for the two methods' ASs. To evaluate the contribution to the alignment quality of each of the two NCFs and two ASs, we mix and match these, resulting in a total of four different combinations. We then use each combination to produce alignments for synthetic networks with known ground

truth node mapping as well as for real-world networks without known ground truth node mapping, and we evaluate the quality of each alignment with respect to five topological and two biological alignment quality measures.

In general, we find that MI-GRAAL's NCF is superior to GHOST's NSF, while the superiority of the methods' ASs is data-dependent. Hence, for those network data on which GHOST's AS is superior to MI-GRAAL's AS, we propose the combination of MI-GRAAL's NCF and GHOST's AS as a new superior network aligner.

While fairly evaluating MI-GRAAL's and GHOST's NCFs and ASs, we approach two additional important research questions that, to our knowledge, have not been asked systematically in the context of network alignment thus far: 1) how much of the node similarity information within the NCF should come from protein sequence data compared to network topology data, and 2) how large the size of the neighborhoods of the compared nodes from different networks should be when generating topological similarity information within the NCF. Current GNA methods generally use a seemingly arbitrary amount of sequence information in their NCF, and also, they assume that the larger the size of a node's neighborhood, the better the alignment quality. Thus, in this study, we evaluate whether these "state-of-the-art" choices are actually appropriate.

In general, we find that which amount of sequence information is used within NCF does not drastically affect neither topological or biological alignment quality, while the effect of topological information is drastic. Namely, using no topological information within NCF results in poor topological and sometimes even biological alignment quality. Hence, topology takes precedence over sequence when it comes to improving alignment quality. Also, we find that using larger network neighborhood sizes within NCF in most cases leads to better alignment quality than using smaller neighborhood sizes. However, it is not always the case that the largest neighborhood size is the best; in many cases, the second largest size is the best. Therefore, using this size instead of the largest one would drastically decrease computational complexity of the given method without decreasing its accuracy.

2 Methods

2.1 Data sets

We use two popular benchmark sets of networks in this study: 1) synthetic networks with known ground truth node mapping and 2) real-world protein-protein interaction (PPI) networks without known ground truth node mapping [19, 32, 34, 28, 9].

The synthetic network data with known node mapping consists of a high-confidence yeast PPI network, which has 1,004 proteins and 8,323 PPIs [7, 18, 26, 19, 32, 34], and five additional networks that add noise to the yeast network. Noise is the addition to the yeast network of low-confidence edges from the same data set [7], and each of the five additional noisy networks adds $x\%$ noise to the original network, where x varies from 5% to 25% in increments of 5%. In this network set, we align the original yeast network to each of the synthetic networks with $x\%$ noise, resulting in the total of five network pairs to be aligned.

The real-world PPI network data without known node mapping consists of PPI networks of the following four species: *S. cerevisiae* (yeast/Y), *D. melanogaster* (fly/F), *C. elegans* (worm/W), and *H. sapiens* (human/H). The yeast, fly, worm, and human networks have 3,321 proteins and 8,021 PPIs, 7,111 proteins and 23,376 PPIs, 2,582 proteins and 4,322 PPIs, and 6,167 proteins and 15,940 PPIs, respectively [4]. In this network set, we align PPI networks for each pair of species, resulting in the total of six network pairs to be aligned.

We note that the synthetic network data is not truly synthetic, as both the original yeast network and the noise in terms of the lower-confidence PPIs come from an actual experimental

study [7]. We refer to this network set as synthetic simply because we know the known ground truth node mapping, unlike for the real-world PPI network set. Also, we note that the synthetic network data encompasses "co-complex" PPIs obtained by affinity purification followed by mass-spectrometry (AP/MS), among other PPI types, while the real-world PPI network data consists of "binary" yeast two-hybrid (Y2H) PPIs. Another difference between the two network sets is for the synthetic data the smaller (original yeast) network is an exact subgraph of the larger (noisy) network, whereas this is not the case for networks of different species in the real-world data.

For evaluating the amount of sequence that should be used within NCF when generating an alignment, we use protein sequence similarity data. This data comes from BLAST bit-values from the NCBI database [1].

When evaluating the biological alignment quality with respect to functional enrichment of the aligned nodes, we use Gene Ontology (GO) annotation data from our recent study [28, 9], to allow for fair and consistent method evaluation.

2.2 Existing network aligners and their NCFs and ASs

MI-GRAAL's NCF. MI-GRAAL improves upon its predecessors, GRAAL [18] and H-GRAAL [26], by using the same NCF (see below) but by combining GRAAL's and H-GRAAL's ASs into a new superior AS (see below).

MI-GRAAL's NCF relies on the concept of small induced subgraphs called *graphlets* (Fig. 1) [27, 23, 24, 25, 39, 14]. All 2-5-node graphlets are considered. Because of the small-world nature of real-world networks, using larger graphlets would unnecessarily increase the computational complexity needed the count the graphlets [18, 26]. Based on the graphlets, the *node graphlet degree vector* (node-GDV) is computed for each node in each network, which counts how many times the given node touches each of the 2-5-node graphlets, i.e., each of their 73 node symmetry groups (or *automorphism orbits*; Fig. 1). As such, node-GDV captures up to a 4-deep network neighborhood of the node of interest. By comparing node-GDVs of two nodes to compute their *node-GDV-similarity*, and by doing so between each pair of nodes in different networks, one is able to capture pairwise topological node similarities between the different networks. MI-GRAAL also allows for integration of other node similarity measures into its NCF, such as protein sequence similarity.

MI-GRAAL's AS. GRAAL's AS utilizes a seed-and-extend approach to greedily maximize the total NCF over all aligned nodes. H-GRAAL, on the other hand, finds optimal alignments with respect to the total NCF by using the Hungarian algorithm to solve the linear assignment problem. MI-GRAAL's AS combines GRAAL's greedy seed-and-extend approach with H-GRAAL's optimal AS into a superior AS. For details on MI-GRAAL's AS, see the original publication [19].

GHOST's NCF. GHOST's NCF takes into account a node's k-hop neighborhood ($k = 4$), which is the induced subgraph on all nodes whose shortest path distance from the node in question is less than or equal to k (Fig. 2). Intuitively, GHOST's NCF computes topological distance (or equivalently similarity) between two nodes from different networks by comparing the nodes' "spectral signatures", where the spectral signature of a node is based on subgraph counts in the node's the k-hop neighborhood [32]. For more details on GHOST's NCF, see the original publication [32]. In our study, we consider $k = 1, 2, 3, 4$, which allows for a fair comparison of GHOST's NCF to MI-GRAAL's NCF when varying the size of network neighborhood that is considered within the NCFs (Section 2.3).

GHOST's AS. Unlike MI-GRAAL's AS that solves linear assignment problem, GHOST's AS solves the quadratic assignment problem (Fig. 3 illustrates this). For further details on GHOST's

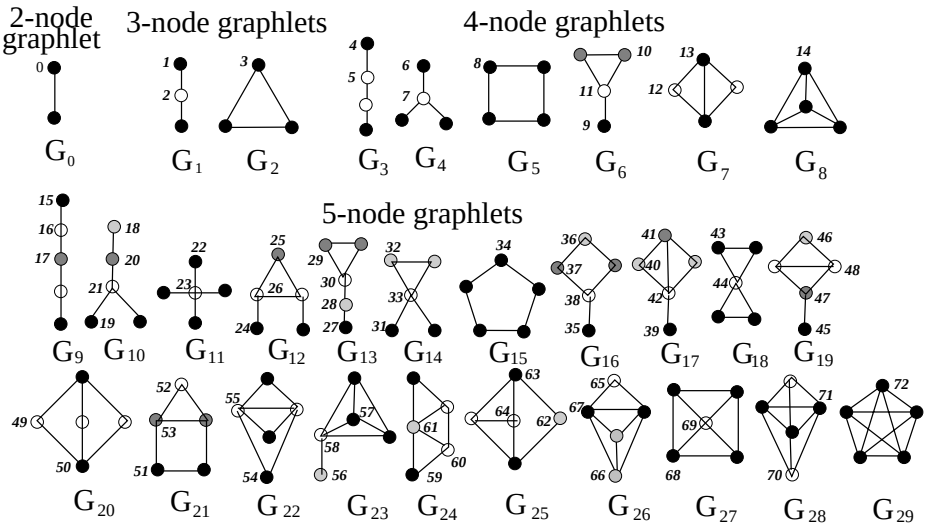

Figure 1: Illustration of MI-GRAAL's NCF, which uses the thirty 2-, 3-, 4-, and 5-node graphlets $G_0, G_1, ..., G_{29}$ and their "node symmetry groups", also called automorphism orbits, numbered 0, 1, 2, ..., 72. In a graphlet G_i, $i \in \{0, 1, ..., 29\}$, nodes belonging to the same orbit are of the same shade. For details, see the original publication [33].

AS, refer to the original publication [32].

2.3 Aligners resulting from combining existing NCFs and ASs, and their parameters

Mixing and matching different NCFs and ASs. To fairly evaluate the two algorithmic components of MI-GRAAL and GHOST, we aim to first compare the two NCFs under the same AS, for each of the two ASs. We then aim to compare the two ASs under the same NCF, for each of the two NCFs. This results in a total of four aligners, i.e., different combinations of the two methods' NCFs and ASs. However, GHOST does not allow the user to import their own (e.g., MI-GRAAL's) NCF into its AS, so we are unable to study the combination of MI-GRAAL's NCF and GHOST's AS. Thus, in total, we consider three different aligners (Table 1).

Aligner	Node Cost Function	Alignment Strategy
M-M	MI-GRAAL	MI-GRAAL
G-M	GHOST	MI-GRAAL
G-G	GHOST	GHOST

Table 1: The three aligners considered in this study. The first letter in the aligner represents NCF of the aligner, while the second letter represents AS of the aligner.

Varying the amount of sequence versus topological information within NCF. An additional goal of this paper is to determine the most appropriate amount of sequence information versus topological information to be included into NCF. Thus, for each aligner, we generate NCFs with varying amounts of sequence and topology information, as $\alpha T + (1-\alpha)S$, where T represents topological similarity score (e.g. node-GDV-similarity) and S represents sequence similarity score. We vary α from 0 to 1 in increments of 0.1.

Varying the size of network neighborhood within NCF. Further, we aim to determine the most appropriate neighborhood size that should be used within NCF when producing an alignment.

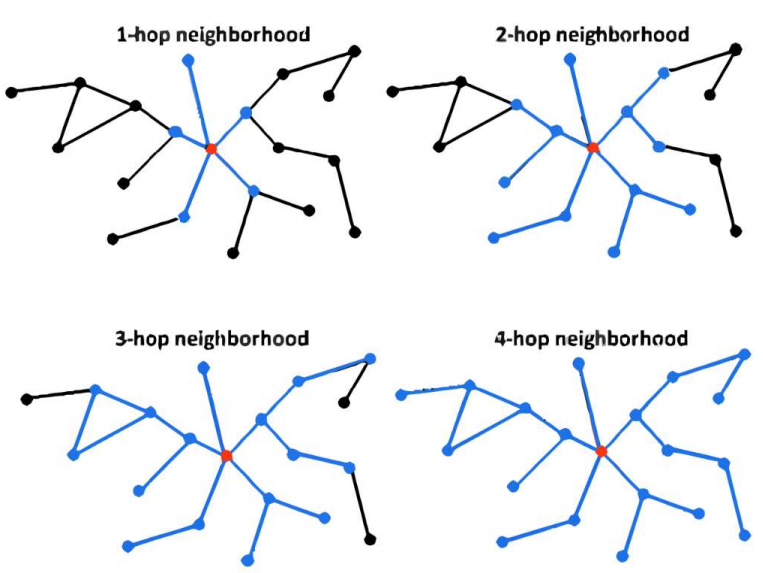

Figure 2: Illustration of GHOST's NCF, which compares two nodes in different networks with respect to similarity of each of their k-hop neighborhoods, $k = 1, 2, 3, 4$. All blue edges and blue nodes are within the given k-hop neighborhood of the red node.

Thus, for each aligner (and for each value of α), we also consider four different neighborhood sizes, as described in Table 2. We note that although we have tried to classify under the same neighborhood size label (e.g. T1 in Table 2) graphlet sizes considered within MI-GRAAL's NCF and k-hop values considered within GHOST's NCF, we note that it is not necessarily the case that the neighborhood of a node that is covered by graphlets of a given size and the neighborhood of the same node that is covered by the corresponding k-hop value match exactly. That is, for example, 2-3-node graphlets and 2-hop neighborhood (both corresponding to T2 in Table 2) do not necessarily cover exactly the same amount of network topology. Yet, we have aimed to provide as accurate as possible classification in Table 2, in order to allow for as fair as possible comparison of the two methods' NCFs under varying sizes of network neighborhoods.

Neighborhood size	Graphlet size (used by MI-GRAAL's NCF)	k-hop neighborhood (used by GHOST's NCF)
T1	2-node graphlets	1-hop neighborhood
T2	2-3-node graphlets	2-hop neighborhood
T3	2-4-node graphlets	3-hop neighborhood
T4	2-5-node graphlets	4-hop neighborhood

Table 2: The four neighborhood sizes that we vary within each aligner.

Implementation details. Types of scores that MI-GRAAL and GHOST take in as input are different: MI-GRAAL looks at node similarities (the higher the score, the more similar the nodes), while GHOST looks at node distances (the lower the score, the more similar the nodes). We carefully take this into account to allow for fair method comparison. For example, to ensure that neither NCF has an advantage due to the format of the scores, we normalize all scores. That is, node similarity scores used in MI-GRAAL can exceed the value 1, while no scores generated by GHOST are greater than one. To make the two sets of scores comparable, we scale MI-GRAAL's node similarity scores to the [0-1] range by dividing each of the scores by the maximum similarity score.

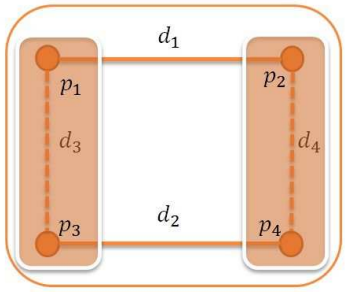

Figure 3: Intuitive comparison of MI-GRAAL's and GHOST's ASs. Let us assume that we are aligning two graphs $G_1(V_1, E_1)$ and $G_2(V_2, E_2)$. Let $p_1, p_3 \in V_1$, let $p_2, p_4 \in V_2$, and let the NCF distance (er equivalently, similarity) between the node pairs be d_1, d_2, d_3, d_4, as illustrated. MI-GRAAL's alignment strategy only considers the values d_1 and d_2 when creating an alignment, while GHOST's AS considers the values d_1,d_2,d_3, and d_4 when doing so.

Then, when we give MI-GRAAL's NCF as input into GHOST's AS, because GHOST deals with distances rather than similarities, we take one minus MI-GRAAL's normalized NCF and then plug in the resulting node scores into GHOST's AS. We do the analogous procedure to convert GHOST's node distance scores into similarity scores when plugging GHOST's NCF into MI-GRAAL's AS, which deals with similarities rather than distances.

Further, MI-GRAAL's NCF returns all pairwise node similarity scores between two networks. However, GHOST's NCF returns only a subset of the lowest of all pairwise node distance scores, depending on the network size. To complete the GHOST's pairwise node score matrix and thus allow for its to be given as input into MI-GRAAL's AS, we assign a score equal to the highest distance score returned by GHOST to all node pairs for which GHOST did not return a distance score.

Finally, current implementation of MI-GRAAL's AS does not function properly when a large pairwise node similarity matrix is plugged into it. Thus, MI-GRAAL's AS has had difficulty aligning two largest networks from our study, those of fly and human. As a solution, we create a matrix that contains only the top 21 million node similarity scores of the original entire node similarity matrix, this being the maximum that our computational resources would process. With this adjustment, we are successfully able to generate all fly-human alignments.

2.4 Network alignment quality measures

We use well established network alignment quality measures [28, 9, 34]. Let $G_1(V_1, E_1)$ and $G_2(V_2, E_2)$ be two graphs such that $|V_1| \leq |V_2|$. An alignment of G_1 to G_2 is a total injective function $f : V_1 \to V_2$; every element of V_1 is matched uniquely with an element of V_2. Let us denote by E_2' the set of edges from G_2 that exist between nodes in G_2 that are aligned by f to nodes in G_1.

Topological evaluation. We use five measures of topological alignment quality:

1. *Node correctness (NC):* If $h : V_1 \to V_2$ is the correct ground truth node mapping between G_1 and G_2 (when such mapping is known), then NC of alignment f is: $NC = \frac{|\{u \in V_1 : h(u) = f(u)\}|}{|V_1|} \times 100\%$ [18]. This measure can be computed only for alignments of the synthetic network set with known ground truth node mapping (Section 2.1). All remaining measures (listed below) can be computed for the real network set with unknown node mapping as well.

2. *Edge correctness (EC):* EC is the percentage of edges from G_1, the smaller network (in terms of the number of nodes), which are aligned to edges from G_2, the larger network [18]. Formally, $EC = \frac{|E_1 \cap E_2'|}{|E_1|} \times 100\%$, where the numerator is the number of "conserved" edges, i.e., edges that are aligned under the given node mapping. The larger the EC score, the better the alignment.

3. *Induced conserved structure (ICS)*: $ICS = \frac{|E_1 \cap E_2'|}{|E_2'|} \times 100\%$. EC might fail to differentiate between alignments that one might intuitively consider to be of different topological quality [32], since it is defined with respect to edges in E_1. For example, aligning a k-node cycle in G_1 to a k-node cycle in G_2 would result in the same EC as aligning a k-node cycle in G_1 to a k-node clique (complete graph) in G_2. Clearly, the former is intuitively a better alignment than the latter, since no edges that exist between the k nodes in G_2 are left unaligned in the first case, whereas many edges are left unaligned in the second case. Since ICS is defined with respect to edges in E_2', it would have the maximum value of 100% when aligning a k-node cycle to a k-node cycle, and it would have a lower value when aligning a k-node cycle to a k-node clique [34]. The larger the ICS, the better.

4. *Symmetric substructure score (S^3)*: EC penalizes the alignment for having misaligned edges in the smaller network. ICS penalizes the alignment for having misaligned edges in the larger network. S^3 on the other hand, aims to improve upon EC and ICS by penalizing for misaligned edges in both the smaller and larger network. $S^3 = \frac{|E_1 \cap E_2'|}{|E_1| + |E_2'| - |E_1 \cap E_2'|} \times 100\%$. For details, see the original publication [34].

5. The size of the *largest connected common subgraph (LCCS)* [18], which we use for the following reason. Of two alignments with similar EC, ICS, or S^3 scores, one could expose large, contiguous, and topologically complex regions of network similarity, while the other could fail to do so. Thus, in addition to counting aligned edges or nodes that participate in the aligned edges, it is important that the aligned edges cluster together to form large connected subgraphs rather than being isolated. Hence, we define a connected common subgraph (CCS) as a connected subgraph (not necessarily induced) that appears in both networks [26]. We measure the size of the largest CCS (LCCS) in terms of the number of nodes as well as edges. Namely, we compute the LCCS score as in our recent work [34]. First, we count N, the percentage of nodes from G_1 that are in the LCCS. Then, we count E, the percentage of edges that are in the LCCS out of all edges that could have been aligned between the nodes in the LCCS. That is, E is the minimum of the number of edges in the subgraph of G_1 that is induced on the nodes from the LCCS, and the number of edges in the subgraph of G_2 that is induced on the nodes from the LCCS [34]. Finally, we compute their geometric mean as $\sqrt{(N \times E)}$, in order to penalize alignments that have small N or small E. Large values of this final LCCS score are desirable.

Biological evaluation. Only alignments in which many aligned node pairs perform the same function should be used to transfer function from annotated parts of one network to unannotated parts of another network [34]. Hence, we measure Gene Ontology (GO) [40] enrichment of aligned proteins pairs, i.e., the percentage of protein pairs in which the two proteins *share* at least one GO term, out of all aligned protein pairs in which both proteins are annotated with at least one GO term. We refer to this percentage as *GO correctness (GO)*. We do this with respect to complete GO annotation data, independent of GO evidence code. Also, since many GO annotations have been obtained via sequence comparison, and since some of the aligners use sequence information, we repeat the analysis considering only GO annotations with experimental evidence codes, in order to

avoid the circular argument. In this case, we refer to GO correctness as *experimental GO correctness (EXP)*. The higher the GO and EXP values, the better [34].

3 Results and discussion

We aim to answer the following three main questions in the context of network alignment: **1)** What is the best NCF and the best AS, and is there perhaps a combination of one existing method's NCF and another existing method's AS that is the superior aligner (Section 3.1)? **2)** How much sequence versus topological information to use within NCF (Section 3.2)? **3)** How large the size of network neighborhoods of compared nodes to consider within NCF (Section 3.3)? In addition, we comment on relationships between different alignment quality measures (Section 3.4). Finally, we conclude in Section 4.

3.1 What is the best NCF and the best AS?

By comparing M-M and G-M aligners, we can fairly compare the two NCFs under the same (MI-GRAAL's) AS. Also, by comparing G-M and G-G, we can fairly compare the two ASs under the same (GHOST's) NCF. See Section 2.3 for details.

3.1.1 Synthetic networks with known node mapping

Overall, GHOST's NCF is slightly superior to that of MI-GRAAL (Fig. 4 (a)-(b)). Also, GHOST's AS is superior to MI-GRAAL's AS (Fig. 4 (a)-(b)). However, these findings are based on *all* alignments (with known node mapping) for all values of α, all neighborhood sizes, and all measures of alignment quality combined (Section 2.3), which might not be fair. Thus, in Fig. 5 (a)-(c), for each aligner, for each alignment quality measure, we show results for the *best* alignments over all values of α and all neighborhood sizes, for three out of all five network pairs (for the remaining network pairs, see the Supplement). Now, the general trend (and especially with respect to NC as the most accurate ground truth measure of alignment quality) is that the best scores for M-M are either comparable or superior to those of G-M, indicating slight superiority of MI-GRAAL's NCF over GHOST's. Nonetheless, G-G still always outperforms G-M, indicating superiority of GHOST's AS over MI-GRAAL's AS.

It is possible to break down the above results and study how the ranking of the different NCFs and ASs changes with the change in the value of α, which corresponds to the amount of topological similarity information used within NCF (see the Supplement). In general, MI-GRAAL's NCF is comparable to GHOST's NCF across all α values, as M-M and G-M scores are similar. On the other hand, GHOST's AS shows superiority over MI-GRAAL's AS, as G-G consistently results in higher scores than G-M. We note that we show that the value of α does not greatly affect alignment quality (Section 3.2).

It is also possible to break down the above results even further and study how the ranking of the different NCFs and ASs changes with the change in the neighborhood size that is considered within NCF (see the Supplement). In general, for the smaller neighborhood sizes (T1 and T2), GHOST's NCF generally produces comparable or superior results to MI-GRAAL's NCF, as G-M scores are higher than M-M scores. However, for the larger neighborhood sizes (T3 and T4), MI-GRAAL's NCF is comparable or superior to GHOST's NCF. And because we show that the larger neighborhood sizes (T3 and T4) are overall superior (Section 3.3), this means that overall MI-GRAAL's NCF is comparable to or superior to GHOST's NCF. On the other hand, in general, for all network sizes, GHOST's AS consistently outperforms MI-GRAAL's AS, as G-G scores is typically higher than G-M scores.

3.1.2 Real networks with unknown node mapping

Overall, unlike for the synthetic network data set with known node mapping, on the real network data set with unknown mapping, MI-GRAAL's NCF is now comparable or superior to that of GHOST (Fig. 4 (c)-(d)). Further, MI-GRAAL's AS is now comparable or superior to GHOST's AS (Fig. 4 (c)-(d)). We confirm these findings even when we limit from *all* alignments (Fig. 4 (c)-(d)) to the *best* alignments only (just as above) (Fig. 5 (d)-(f)) (see the Supplement).

When zooming into the results further to observe the effect of the α parameter, in general, for all values of α, MI-GRAAL's NCF is comparable or superior to GHOST's NCF and MI-GRAAL's AS is comparable to GHOST AS across all values of α (see the Supplement).

The same holds independent on the neighborhood size that is considered within NCF (see the Supplement).

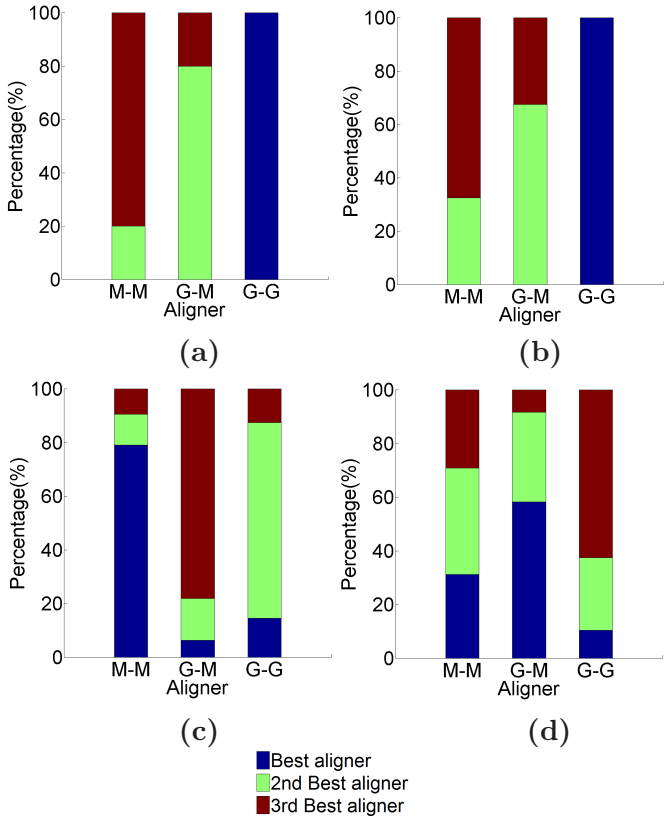

Figure 4: The ranking of the three aligners (M-M, G-M, and G-G) over *all* alignments for all values of α and all neighborhood sizes, with respect to: **(a)** all topological scores of all alignments with known ground truth node mapping, **(b)** all biological scores of alignments with known node mapping, **(c)** all topological scores of alignments with unknown node mapping, and **(d)** all biological scores of alignments with unknown node mapping.

3.1.3 Summary

Which NCF or AS is the best *overall* is not easy to determine, as the results are data-dependent. But when we limit analyses of each aligner to the *best* alignments over all parameters, M-M is comparable or superior to G-M, indicating that MI-GRAAL's NCF is better than GHOST's NCF,

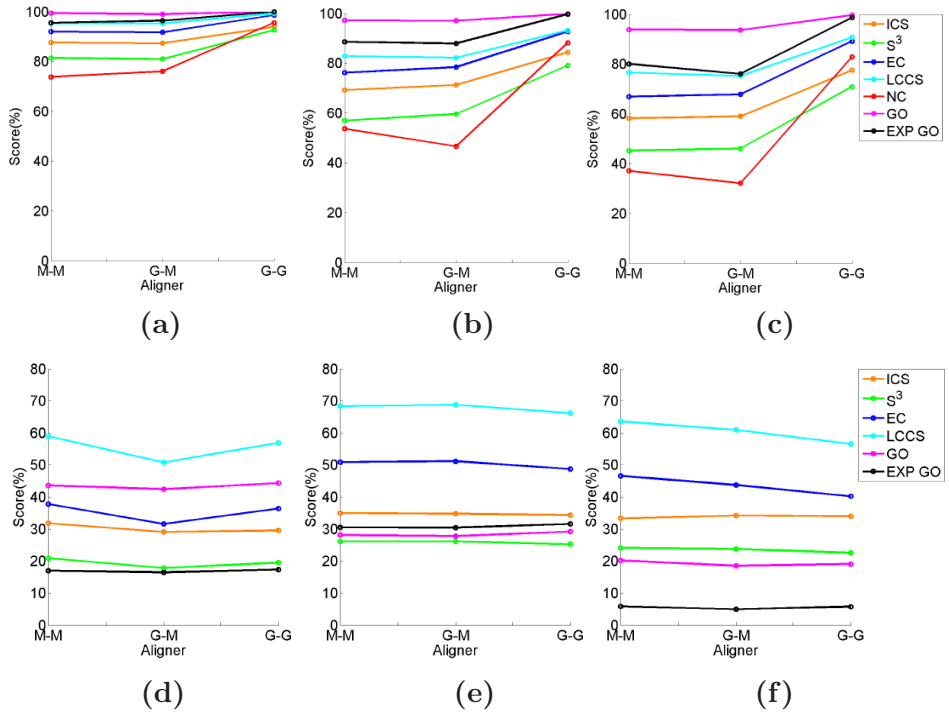

Figure 5: Alignment quality results of the three aligners (M-M, G-M, and G-G) for *best* alignments over all values of α and all neighborhood sizes, for (a)-(c) three network pairs with known node mapping (yeast-yeast 5%, yeast-yeast 10%, and yeast-yeast 15%, respectively) and (d)-(e) three network pairs with unknown mapping (human-yeast, human-worm, and worm-yeast, respectively). For equivalent results for the remaining network pairs, see the Supplement.

while the performance of G-M versus G-G, i.e., of MI-GRAAL's AS versus GHOST's AS, is still data-dependent.

Therefore, the graphlet-based measure of topological node similarity [27] that MI-GRAAL uses (along with many other network aligners [18, 26, 28, 9] or even network clustering methods [27, 24, 39]) remains the state-of-the-art, even when compared to the spectral signature-based node similarity measure that GHOST uses (and especially compared to PageRank-based node similarity measure that aligners from the IsoRank family use, as we already showed in our recent study [28, 9]).

Our results indicate that the slight superiority of GHOST (i.e., G-G) over MI-GRAAL (i.e., M-M) that was claimed in the original GHOST publication [32] seems to come from GHOST's AS and *not* its NCF, which is not surprising, since GHOST's AS solves quadratic assignment problem whereas MI-GRAAL's AS deals only with linear assignment problem. Further, our results indicate that the combination of MI-GRAAL's NCF and GHOST's AS (i.e., M-G) could be a new aligner that is superior to the existing MI-GRAAL (i.e., M-M) and GHOST (ie., G-G) aligners on at least some data sets. Unfortunately, explicitly testing this is not possible with the current implementation of GHOST, as per our conversation with the authors of GHOST, the current implementation is too complex to modify to allow for plugging MI-GRAAL's (or any other method's) NCF into GHOST's AS.

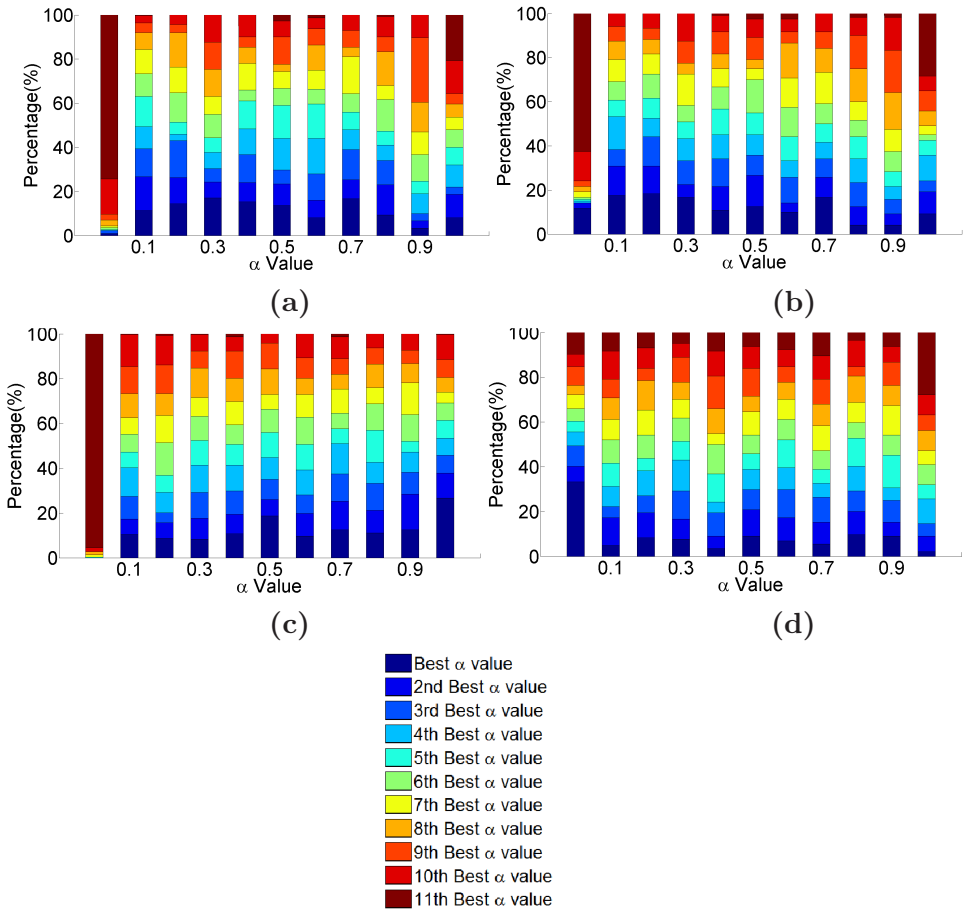

Figure 6: The ranking of the 11 values of α (from 0 to 1 in increments in 0.1) over *all* alignments for all aligners and all neighborhood sizes, with respect to: **(a)** all topological scores of alignments with known ground truth node mapping, **(b)** all biological scores of alignments with known node mapping, **(c)** all topological scores of alignments with unknown node mapping, and **(d)** all biological scores of alignments with unknown node mapping.

3.2 The amount of sequence versus topological information within NCF?

Recall that we vary the amount of topological node similarity information within NCF with the α parameter (where α of 0 means that no topology information is used, i.e., that only sequence information is used, whereas α of 1 means that only topology information is used; Section 2.3). Here, we study the effect of the α parameter on alignment quality.

3.2.1 Synthetic networks with known node mapping

Overall, the value of α does not affect alignment quality, as long as some amount of topological information is used. That is, only $\alpha = 0.0$ results in completely inferior alignments, especially with respect to topological alignment quality, whereas all other values of alpha are more-less comparable (Fig. 6 (a)-(b)). It is expected that the larger the value of α, i.e., the more of topological information is used within NCF, the better the topological alignment quality. Again, this is exactly what we observe (Fig. 6 (a)). It is also expected that that the smaller the value of α, i.e., the more of sequence information is used within NCF, the better the biological alignment quality. Surprisingly,

this is not what we observe (Fig. 6 (b)): larger values of α (e.g., 0.7) result in more of high-quality alignments than $\alpha = 0$.

When zooming into the results further to observe the effect of the aligner, in general, we see the same trends as above independent of the aligner (see the Supplement). Namely, the results from Fig. 6 (a)-(b) hold independent on which NCF or AS is used. Further, there is no difference in the results across the two NCFs (Fig. 7 (a) and (b)). There is only a minor difference in the results across the two ASs, in the sense that the results are somewhat more stable across different αs for GHOST's AS than for MI-GRAAL's AS (Fig. 7 (b) and (c)). Also, GHOST's AS suggests that in addition to not using $\alpha = 0$ (i.e., sequence alone), one should not use $\alpha = 1$ either (i.e., topology alone); but other than that, the choice of α still has no major effect (Fig. 7 (c)).

When zooming into the results from Fig. 6 (a)-(b) further to observe the effect of the neighborhood size, we see that the results hold independent of the neighborhood size (see the Supplement).

3.2.2 Real networks with unknown node mapping

The results that we observe for the synthetic networks in general hold for this network set as well. Namely, $\alpha = 0$ results in the worst topological alignment quality, while the other α values are somewhat comparable, with a slight dominance of the larger values, as expected (Fig. 6 (c)). Interestingly, for this network set, the lowest value of $\alpha = 0$ results in the most of highest-scoring alignments with respect to biological alignment quality; yet, even the largest αs often lead to good alignments with respect to biological alignment quality (Fig. 6 (d)).

When zooming into the results further to observe the effect of the aligner, as with synthetic networks, the general results from Fig. 6 (c) and (d) hold independent of the aligner for real networks as well (see the Supplement). However, unlike for synthetic networks, for real networks we now see result stability across all NCFs and all ASs, and not just for GHOST's AS. Also, now GHOST's AS no longer suggests that $\alpha = 1$ should not be used.

When zooming into the results from Fig. 6 (c) and (d) further to observe the effect of the neighborhood size, just as with the synthetic networks, we again see that the results hold independent of the neighborhood size (see the Supplement).

3.2.3 Summary

Overall, at least some amount of topological information should be included within NCF, as this results in good topological as well as biological alignment quality. While $\alpha = 0.0$ may (but does not always) result in biologically high-quality alignments, in every case it fails to produce topologically superior results. Thus, $\alpha = 0.0$ should not be used.

3.3 The size of nodes' neighborhoods within NCF?

Intuitively, one would expect that the increase in the size of nodes' network neighborhoods within NCF (i.e., in the amount of network topology) would result in higher-quality alignments. However, this assumption has not been tested to date. Instead, the existing methods blindly use the largest neighborhood size that is allowed by available computational resources (that is, MI-GRAAL uses all 2-5-node graphlets, whereas GHOST uses $k = 4$; Section 2.3). Thus, within each aligner, we vary the neighborhood size from T1 to T4 (Table 2) to systematically evaluate the effect of this parameter.

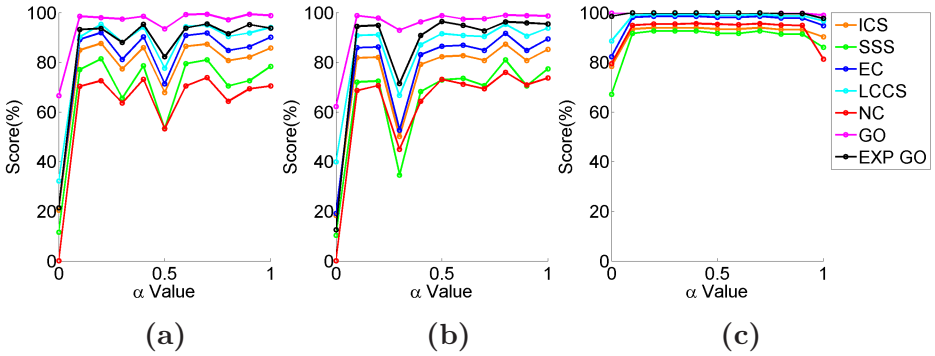

Figure 7: Detailed illustration of the effect of the α parameter for **(a)** M-M, **(b)** G-M, and **(c)** G-G aligners. In particular, results are shown for the yeast-yeast 5% alignment and for the neighborhood size T4. For other network pairs and other neighborhood sizes, see the Supplement for synthetic network data and see the Supplement for real-world PPI network data.

3.3.1 Synthetic networks with known node mapping

Overall, the larger the neighborhood size, the better the alignment quality, even though all neighborhood sizes except T1 can in some cases result in higher-quality alignments than any other neighborhood size (Fig. 8 (a)-(b)). That is, for some values of network alignment parameters, smaller neighborhoods can produce higher-quality alignments than larger neighborhoods, which is a surprising result.

When zooming into the results further to observe the effect of the aligner, the general trends from Fig. 8 (a)-(b) still hold independent of the aligner, but some fluctuations in the results exist (see the Supplement). Namely, M-M generally prefers T3 and T4 neighborhood sizes. G-M prefers T2 in addition to T3 and T4, where T3 or T4 are actually inferior to T2 in some cases, depending on the noise level. G-G performs well on of T1-T4, with a slight preference of T3 or T4, depending on the noise level. See Fig. 9 (a) for an illustration.

When zooming into the results further to observe the effect of the α parameter, general trends from Fig. 8 (a)-(b) are overall the same for all values of α (see the Supplement). The only exception is $\alpha = 0$, which should not be used in the first place (Section 3.2.3).

3.3.2 Real networks with unknown node mapping

Unlike for the synthetic networks, the largest neighborhood size (T4) is now not overly dominant over the smaller network sizes. Specifically, for real network data set, it is T3 that is the most dominant, followed by T4 and T2, which are tied, and followed by T1, which is inferior (Fig. 8 (c) and (d)).

When zooming into the results further to observe the effect of the aligner, we see that each aligner has an interesting behavior (see the Supplement). Namely, M-M's and G-G's preference on the neighborhood size is mainly dictated by the choice of species whose networks are aligned. For G-M, in general, the larger neighborhood sizes are preferred; in some cases, depending on the species, G-M prefers T3 more than other neighborhood sizes. See Fig. 9 (b) for an illustration.

When zooming into the results further to observe the effect of the α parameter, just as for synthetic networks, the results from Fig. 8 (c) and (d) do not drastically change with the change of α value (see the Supplement).

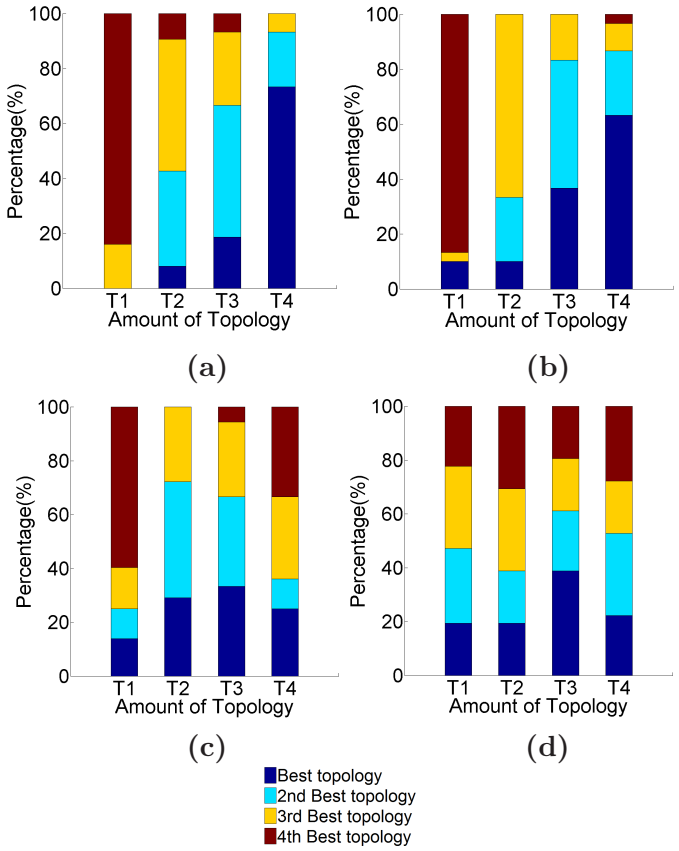

Figure 8: The ranking of the four neighborhood sizes (T1-T4) over *all* alignments for all aligners and all values of α, with respect to: **(a)** all topological scores of alignments with known ground truth node mapping, **(b)** all biological scores of alignments with known node mapping, **(c)** all topological scores of alignments with unknown node mapping, and **(d)** all biological scores of alignments with unknown node mapping.

3.3.3 Summary

In general, the larger the neighborhood size within NCF, the higher the alignment quality. However, it is not necessarily the case that the largest neighborhood size always produces the best alignments nor that it is always dominant to the smaller neighborhood sizes. This means that slightly smaller neighborhood sizes (and T3 in particular) might be desirable, as this could not only produce better alignments in some cases but also decrease the computational complexity of the given method.

3.4 Relationships between different alignment quality measures

We use a total of seven alignment quality measures: the ground truth NC measure that can only be measured in alignments of synthetic networks with known node mapping, four additional topological measures (EC, ICS, S^3, and LCCS), and two biological measures (GO and EXP) (Section 2.4). Here, we briefly comment on the relationship between the different measures.

NC significantly correlates with both topological and biological alignment quality measures (Fig. 10 (a)), which is encouraging. Further, for the synthetic network data set, it is also encouraging that all other measures significantly correlate well (Pearson correlation coefficient of at least 0.8), even though we see some clustering of the topological measures and also of the biological measures

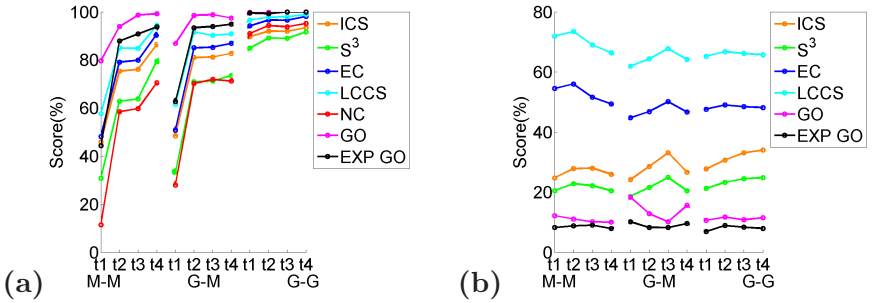

Figure 9: Detailed illustration of the effect of the neighborhood size for **(a)** synthetic and **(b)** real network data. In particular, results are shown for all three aligners, for the yeast-yeast 5% alignment at $\alpha = 0.6$ in panel (a) and for the fly-worm alignment at $\alpha = 0.4$ in panel (b). For other network pairs and other values of α, see the Supplement.

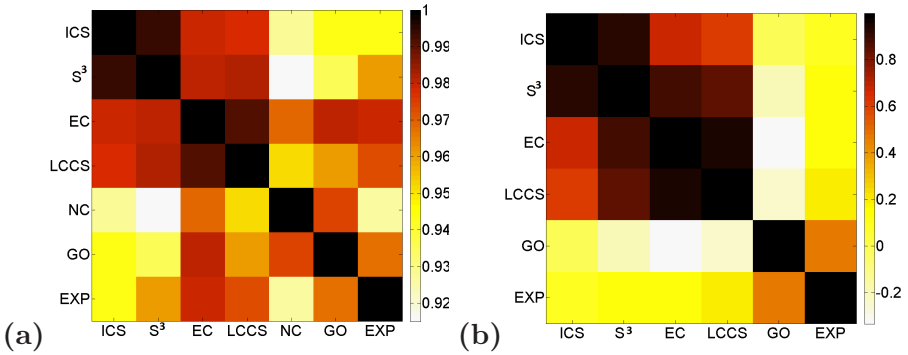

Figure 10: Pairwise correlations between different alignment quality measures for: **(a)** synthetic networks with known ground truth node mapping and **(b)** real networks with unknown node mapping. Correlations were computed over alignments with the highest NC scores in panel (a) and over alignments with the highest EC scores in panel (b) (because we do not known NC scores for alignments of real networks). Note that color scales for the two panels are different.

(Fig. 10 (a)). Interestingly, each of the two biological measures, GO and EXP, correlates better with some of the topological measures (e.g., EC) than with each other.

Unlike for the synthetic network data, for the real network data, the topological measures now correlate poorly with the biological measures (Pearson correlation coefficient of at most 0.2; Fig. 10 (b)). Importantly, this implies that for the real network data set it might be hard to produce an alignment that is of excellent quality both topologically and biologically. Also, while we again see clustering of the topological measures, the two biological measures now correlate weakly (Fig. 10 (b)), indicating that the choice of GO annotation data obtained by experimental evidence code matters (Section 2.4).

The result differences between the synthetic networks and the real networks could be due to differences in their properties (Section 2.1).

4 Conclusions

We have aimed to systematically answer three questions in the context of MI-GRAAL and GHOST network aligners: 1) what is the contribution of each method's NCF and AS to the alignment

quality, 2) how much sequence versus topology information should be used within NCF when generating an alignment, and 3) how large the size of the neighborhoods of the compared nodes from different networks should be. Our results represent a set of general recommendations for a fair evaluation of any global network alignment method, not just MI-GRAAL and GHOST.

Genomic sequence alignment has revolutionized our biomedical understanding. Biological network alignment has already had similar impacts. And given the tremendous amounts of biological network data that continue to be produced, network alignment will only continue to gain importance. The hope is that it could lead to new discoveries about the principles of life, evolution, disease, and therapeutics.

Acknowledgements

We thank Dr. R. Patro and Dr. C. Kingsford for their assistance with running GHOST. This work was supported by the National Science Foundation CCF-1319469 and EAGER CCF-1243295 grants.

References

[1] S. F. Altschul, W. Gish, W. Miller, and D. J. Lipman. Basic local alignment search tool. *Journal of Molecular Biology*, 215:403–410, 1990.

[2] J. Berg and M. Lassig. Local graph alignment and motif search in biological networks. *Proceedings of the National Academy of Sciences*, 101(41):14689–14694, 2004.

[3] J. Berg and M. Lassig. Cross-species analysis of biological networks by Bayesian alignment. *Proceedings of the National Academy of Sciences*, 103(29):10967–10972, 2006.

[4] B. J. Breitkreutz, C. Stark, , T. Reguly, L. Boucher, A. Breitkreutz, M. Livstone, R. Oughtred, D. H. Lackner, J. Bahler, V. Wood, K. Dolinski, and M. Tyers. The BioGRID Interaction Database: 2008 update. *Nucleic Acids Research*, 36:D637–D640, 2008.

[5] Giovanni Ciriello, Marco Mina, Pietro H. Guzzi, Mario Cannataro, and Concettina Guerra. AlignNemo: A Local Network Alignment Method to Integrate Homology and Topology. *PLOS ONE*, 7(6), 2012.

[6] Connor Clark and Jugal Kalita. A comparison of algorithms for the pairwise alignment of biological networks. *Bioinformatics*, 2014. DOI: 10.1093/bioinformatics/btu307.

[7] S.R. Collins, P. Kemmeren, X.C. Zhao, J.F. Greenblatt, F. Spencer, F.C.P. Holstege, J.S. Weissman, and N.J. Krogan. Toward a comprehensive atlas of the phyisical interactome of Saccharomyces cerevisiae. *Molecular Cell Proteomics*, 6(3):439–450, 2007.

[8] Mohammed El-Kebir, Jaap Heringa, and Gunnar W Klau. Lagrangian relaxation applied to sparse global network alignment. In *Pattern Recognition in Bioinformatics*, pages 225–236. Springer, 2011.

[9] F. Faisal, H. Zhao, and T. Milenkovic. Global network alignment in the context of aging. *Computational Biology and Bioinformatics, IEEE/ACM Transactions on Computational Biology and Bioinformatics*, PP(99), 2014.

[10] Fazle E Faisal and Tijana Milenković. Dynamic networks reveal key players in aging. *Bioinformatics*, 30(12):1721–1729, 2014.

[11] J. Flannick, A. Novak, S.S. Balaji, H.M. Harley, and S. Batzglou. Graemlin general and robust alignment of multiple large interaction networks. *Genome Research*, 16(9):1169–1181, 2006.

[12] Jason Flannick, Antal Novak, Chuong B. Do, Balaji S. Srinivasan, and Serafim Batzoglou. Automatic parameter learning for multiple network alignment. In *Proceedings of the 12th annual international conference on Research in computational molecular biology*, pages 214–231, 2008.

[13] X. Guo and A.J. Hartemink. Domain-oriented edge-based alignment of protein interaction networks. *Bioinformatics*, 25(12):i240–1246, 2009.

[14] Y. Hulovatyy, R.W. Solava, and T. Milenković. Revealing missing parts of the interactome via link prediction. *PLOS ONE*, 9(3):e90073, 2014.

[15] Brian P. Kelley, Bingbing Yuan, Fran Lewitter, Roded Sharan, Brent R. Stockwell, and Trey Ideker. PathBLAST: a tool for alignment of protein interaction networks. *Nucleic Acids Research*, 32:83–88, 2004.

[16] G.W. Klau. A new graph-based method for pairwise global network alignment. *BMC Bioinformatics*, 10(Suppl 1):S59, 2009.

[17] M. Koyuturk, Y. Kim, U. Topkara, S. Subramaniam, W. Szpankowski, and A. Grama. Pairwise alignment of protein interaction networks. *Journal of Computational Biology*, 13(2), 2006.

[18] Oleksii Kuchaiev, Tijana Milenković, Vesna Memišević, Wayne Hayes, and Nataša Pržulj. Topological network alignment uncovers biological function and phylogeny. *Journal of The Royal Society Interface*, 7(50):1341–1354, 2010.

[19] Oleksii Kuchaiev and Nataša Pržulj. Integrative network alignment reveals large regions of global network similarity in yeast and human. *Bioinformatics*, 27(10):1390–1396, 2011.

[20] Juanzi Li, Jie Tang, Yi Li, and Qiong Luo. RiMOM: A dynamic multistrategy ontology alignment framework. *Knowledge and Data Engineering, IEEE Transactions on*, 21(8):1218–1232, 2009.

[21] Zhi Liang, Meng Xu, Maikun Teng, and Liwen Niu. NetAlign: A web-based tool for comparison of protein interaction networks. *Bioinformatics*, 22(17):2175–2177, 2006.

[22] C. Liao, K. Lu, M. Baym, R. Singh, and B. Berger. IsoRankN: Spectral methods for global alignment of multiple protein networks. *Bioinformatics*, 25(12):i253–258, 2009.

[23] T. Milenković, J. Lai, and N. Pržulj. GraphCrunch: a tool for large network analyses. *BMC Bioinformatics*, 9(70), 2008.

[24] T. Milenković, V. Memisević, A. K. Ganesan, and N. Pržulj. Systems-level cancer gene identification from protein interaction network topology applied to melanogenesis-related interaction networks. *Journal of the Royal Society Interface*, 7(44):423–437, 2010.

[25] T. Milenković, V. Memišević, A. Bonato, and N. Pržulj. Dominating biological networks. *PLOS ONE*, 6(8):e23016, 2011.

[26] T. Milenković, W.L. Ng, W. Hayes, and N. Pržulj. Optimal network alignment with graphlet degree vectors. *Cancer Informatics*, 9:121–137, 2010.

[27] T. Milenković and N. Pržulj. Uncovering biological network function via graphlet degree signatures. *Cancer Informatics*, 6:257–273, 2008.

[28] Tijana Milenković, Han Zhao, and Fazle E. Faisal. Global network alignment in the context of aging. In *Proceedings of the International Conference on Bioinformatics, Computational Biology and Biomedical Informatics*, BCB'13, pages 23:23–23:32. ACM, 2013.

[29] Marco Mina and Pietro Hiram Guzzi. Improving the robustness of local network alignment: design and extensive assessment of a markov clustering-based approach. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 99(PrePrints):1, 2014.

[30] Arvind Narayanan, Elaine Shi, and Benjamin IP Rubinstein. Link prediction by de-anonymization: How we won the Kaggle social network challenge. In *Neural Networks (IJCNN), The 2011 International Joint Conference on*, pages 1825–1834. IEEE, 2011.

[31] Behnam Neyshabur, Ahmadreza Khadem, Somaye Hashemifar, and Seyed Shahriar Arab. NETAL: a new graph-based method for global alignment of protein-protein interaction networks. *Bioinformatics*, 29(13):1654–1662, 2013.

[32] R. Patro and C. Kingsford. Global network alignment using multiscale spectral signatures. *Bioinformatics*, 28(23):3105–3114, 2012.

[33] N. Pržulj. Biological network comparison using graphlet degree distribution. *Bioinformatics*, 23:e177–e183, 2007.

[34] V. Saraph and T. Milenković. MAGNA: Maximizing Accuracy in Global Network Alignment. *Bioinformatics*, 2014. DOI: 10.1093/bioinformatics/btu409.

[35] R. Sharan and T. Ideker. Modeling cellular machinery through biological network comparison. *Nature Biotechnology*, 24(4):427–433, 2006.

[36] R. Sharan, S. Suthram, R.M. Kelley, T. Kuhn, S. McCuine, P. Uetz, T. Sittler, R.M. Karp, and T. Ideker. Conserved patterns of protein interaction in multiple species. *Proceedings of the National Academy of Sciences*, 102(6):1974–1979, 2005.

[37] R. Singh, J. Xu, and B. Berger. Global alignment of multiple protein interaction networks. *In Proceedings of Pacific Symposium on Biocomputing 13*, pages 303–314, 2008.

[38] Rohit Singh, Jinbo Xu, and Bonnie Berger. Pairwise global alignment of protein interaction networks by matching neighborhood topology. pages 16–31. Springer, 2007.

[39] R.W. Solava, R.P. Michaels, and T. Milenković. Graphlet-based edge clustering reveals pathogen-interacting proteins. *Bioinformatics*, 18(28):i480–i486, 2012.

[40] The Gene Ontology Consortium. Gene Ontology: tool for the unification of biology. *Nature Genetics*, 25:25–29, 2000.

[41] D.B. West. *Introduction to Graph Theory*. Prentice Hall, Upper Saddle River, NJ, 2^{nd} edition, 2001.

[42] M. Zaslavskiy, F. Bach, and J. P. Vert. Global alignment of protein-protein interaction networks by graph matching methods. *Bioinformatics*, 25(12):i259–i267, 2009.